nVy eight-
year
ds who
master
Interne
t. T
hey

V in
C e
to n G .
r f

Pioneer and father of the Internet. Co-inventor of TCP/IP.

t twenty
years
before me.
WE HAD

TO NVENT IT for m to be able to hat they do.”

It would be impossible to understand the history of the Internet without his contribution. This document no doubt reflects why Cerf is affection- ately nicknamed the “father of the Internet,” although he insists on call- ing himself an “Internet pioneer.” Vint is a man who has devoted his life on solving complicated technological problems with simple ideas and brilliant solutions.

Interview conducted[1] on December 8th, 2007 in McLean, Virginia

Cerf was born on June 23rd, 1943, in New Haven, Connecticut, at Yale University Hospital. His mother lived there with him while his father was in the European theater of World War II. He was in the Navy and served as an artillery officer aboard a submarine. Before the atomic bomb was dropped in Japan, he was preparing to be shipped to the Pacific theater, where they had attempted to invade Japan. But after the launch of the nuclear devices, the war ended and they moved from New Haven to Los Angeles in 1946.

So, Vinton Cerf actually grew up in Los Angeles. He lived in the San Fernando Valley, north of the city, in a small town called Van Nuys. He went to high school there, where it turned out that Steve Crocker, Jon Postel and Karl Auerbach had also studied. They weren't however in the same year. Jon Postel was one year behind Steve, and Steve one year behind Vinton Cerf. Steve was his best friend in high school, and he still is, but neither of the two knew Jon Postel back then. In fact, they didn't meet Jon until they finished their degree at UCLA.

After graduating from high school, he enrolled at Stanford University where, from 1961 to 1965, he majored in mathematics and minored in German[2].

Listen to the interview

In 1965, he began working for IBM at the headquarters known as the Los Angeles data center.

In 1967, after studying a PhD in Computer Engineering at UCLA, he found himself working at ARPAnet, developing the first protocol between mainframes, and in 1972 he returned to Stanford as a teacher.

In 1976 he returned to ARPA, and stayed there until 1982 when he joined MCI. In 1986 he passed on to CNRI, to return to MCI in 1994 until 2005.

www.ucla.edu *www.ibm.com* *www.stanford.edu*
www.arpa.mil *www.mci.net* *www.google.com*

Vinton G. Cerf

Interview conducted through discussions June 25th, 1999 (in San José, California)

, May 20th, 2000 (Tarragona), June 12th, 2001 (Stockholm), and completed via electronic correspondence on June 23rd, 2001 (Tenerife), and November 23rd, 2001 (Barcelona). Plus the recorded interview in McLean (Virginia) during the evening of December 8th, 2007.

2 I have been studying German since I was about thirteen years old. My father contracted a German tutor to give me private lessons at home once a week. I went to Stanford in Germany, a campus in an area near Stuttgart, where we lived in a small town called Beutelsbach which had about 3,000 inhabitants and was about 30 km from Stuttgart. The house we lived in was a farm on top of a hill called Landgut Burg, and there were about 70 students there. So Stanford brought in professors and classes were held there, sometimes in English but mostly in German. We got to know all the local people and traveled around Europe during the six months I was there. To Bonn, Prague, Rome, Stockholm, etc., because we made the most of the weekends and, in Europe, the distances are quite short. I came back in December 1962 to continue my studies at Stanford.

Do you remember when you first came into contact with a computer?

It was in 1958, I don't recall what month of the year. It was in Santa Monica, California, at a company called Systems Development Corporation, which was a RAND spin-off that at the time was conducting research classified for the US Air Force. In particular, Paul Baran at RAND Corporation wrote one of the first lengthy papers on what would become packet switching, although he did not call it that at the time. RAND had been working on command and control for the national military and one of the ideas that came out of the RAND work was to put radars in the northern part of Canada looking towards Russia and to transmit the radar information by land line all the way to Santa Monica, California, where the data would be absorbed into a computing complex which was called the Semi Automated Ground Environment (SAGE). This system was built using radio valves. The idea was that the data would be transmitted from the radar and displayed through large radar screens that would be processed by SAGE. There was a huge sensation of being at the center of the system—you literally had to enter the computer to use it. There were rooms full of vacuum valves along the walls. I remember becoming completely mesmerized by the idea that you could have a computer dealing with things from such a long distance away. About two years later, Steve Crocker obtained permission to use some of UCLA's computers. This would have been around 1960. We were taking different courses from each other at Van Nuys High School. Steve and I were best friends and he invited me to go with him to try out the Bendix G-15 computer at UCLA. A little later, I went to Stanford as a student and, of course, immediately took all the courses I could on Computer Engineering. We used a Burroughs B5000 machine that had been upgraded to B5500, and we programmed in ALGOL, or what was called Burroughs ALGOL. So that was very exciting time for me, both at Stanford as undergraduate but also as high school student using computers at UCLA. I have these very

What was your first contact/experience with Internet or the ARPAnet?

My first contact with ARPAnet was with **Len Kleinrock** at the Network Measurement Center[3], which would have been around 1969 or maybe 1968. That was the year the RFQ was issued for the IMPS (Interface Message Processors — the packet switches of the ARPAnet) and interestingly enough Steve Crocker and I bid on the ARPAnet project while we were still graduate students but we were working as consultants for a company called Jacobi Systems at the time. Jacobi submitted a proposal by ARPAnet; he was a competitor of Bolt, Beranek and Newman (BBN) and we went all the way to the last rounds. I think we finished in the last four. I wrote a simulation of the packet-switched network using the GPSS programming language. Jacobi Systems didn't get the contract, but luckily Steve and I were still students and working in Len Kleinrock's group. When BBN won, UCLA got the network measurement center contract and we got to work on that. So, I started working at ARPAnet even though Jacobi Systems didn't win the contract. Jacobi might have been the name of the founder of the company; they were based in Santa Monica.

"In the early days, I worked on this project very passionately because I wanted to solve the Defense Department's problem: they didn't want a proprietary network, and they didn't want to limit themselves to a single network technology."

On the Internet side, it's already obvious because I started working on the project in 1973, so that was when I had my first contact. When we started the project, we called it internetworking" (network interconnection operation), and the idea was to select packet-switched networks

vivid recollections of getting caught up in computers and getting completely overrun with the excitement of using those machines.

3 En el original inglés el centro es denominado: Network Measurement Center.

and connect them to each other. But I remember seeing references to the shortened phrase "Internet" in December 1974 or even earlier. The first specification of the TCP protocol referred to the TCP Internet protocol, and I think we had already adopted the idea of this multiple network, formed by something called the Internet, in normal speech. I always wrote it with a capital 'I'. When the network was extended in 1983 and made accessible, people said: "Well, sometimes you can build pieces of network using Internet technology, which is not part of the public network." So then we thought about writing the networks that used TCP/ IP protocols in lowercase. The Internet began to be deliberately written with a capital 'I' to refer to the ARPA project, while in lowercase letters it referred to networks that used Internet technology but were not connected to the DARPA Internet. In the end, when the Internet became public, I got into the habit of referring to the public Internet in capital letters and any other private network that used the same technology in lowercase letters. But the term Internet was intended to reflect the idea that there were multiple networks connecting to each other. So, the interconnections of the network and the structure in which it was located were the focus of the research.

What has been your professional experience in the development of the Internet?

When I graduated from Stanford in June 1965, I went to work for IBM in what was called the Los Angeles Data Center. It used a time-sharing system for IBM called Quiktran. It was an interactive FORTRAN system that remotely operated the IBM 7044 mainframe computer to write and run FORTRAN programs. I did stress analysis on the buildings, and we had a client who used the service to program algorithms for buying and selling stocks. There was just a wide range of applications, anything you could program in FORTRAN. I ran this operational center as a systems engineer for two years, until mid-1967. Then I got a license to go back to class
because I felt the need to get advanced training in processing.

multi-process synchronization to Jerry, about several computers running in parallel. But I also worked for Len Kleinrock as a senior programmer for the Network Measurement Center. Our job was to inject traffic and collect data from ARPAnet to validate the queue model that **Kleinrock** and his students were developing for packet-switched networks. As a result of these experiences, I was quickly drawn to ARPAnet's activities in the late 1960s, and ended up working with **Steve Crocker, Stephen Carr** and **Jon Postel** in the first protocol between central computers *(host-to-host)* known as NCP (Network Control Protocol). We did this and then presented the system in public in October 1972 at the Hilton Hotel in Washington at the first International Computer Communications Conference (ICCC). Bob Kahn organized the demonstration at the request of Larry Roberts who led the ARPAnet program for ARPA and quite a few of the people that you have already interviewed were at that meeting including Louis Pouzin (Cyclades, France). At that meeting we formed a group called the INWG[4] (the International Network Working Group) which eventually became the IFIP (International Federation for Information Processing) working group 6.1, dealing with data communications. This group of people was quite interested in the notions of computer networking and packet switching in particular. At **Steve Crocker's** recommendation, I became chairman. Although I don't think any specific developments came out of this group as a working group, there were very useful discussions on network architecture and protocols.

Bob Kahn and I presented to that group in September 1973, the basic design of the Internet in a document that we called *INWG # 39*. The group produced a large number of working papers, and some of the people who were interviewed were involved with this group or documented it, acting as editors of the series of notes. I don't know what happened to those notes[5], and it would be interesting to know whether or not they are in the Computer History Museum[6].

At UCLA I did a Master's degree and then a PhD in Computer Engineering for five years, and my thesis tutor was **Gerald Estrin.** I submitted a thesis on

4 International Network Working Group.
5 Alexander McKenzie, a former BBN employee, may know what happened to them and I think Peter Sevcik might also know because I think they were both editors of the series.
6 Known by its initials: CHM. It is located in Mountain View, California.

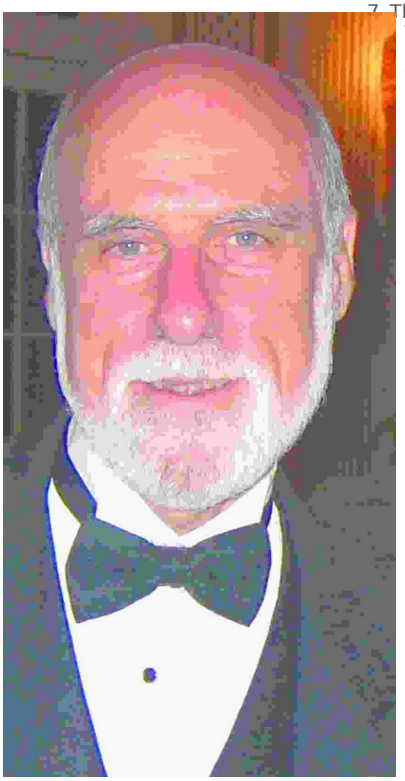

After the demonstration in October 1972, Bob Kahn left BBN (Bolt, Beranek and Newman) and went to ARPA, and I left UCLA and went to Stanford, this time as a faculty member, and joined both the computer science and the electrical engineering departments. There, research on the Internet was initiated and supported by DARPA, and many of the people you interviewed worked for Bob while I was teaching at Stanford.

TCP/IP protocols were developed from 1973 to 1976. The IP part wasn't actually split off until 1977, and it was thanks to the insistence of **Danny Cohen, David Reed** and **Jon Postel** that we separated the IP protocol layer from the TCP layer. It was primarily motivated by the desire to carry real time traffic that didn't necessarily have to all be delivered, whereas TCP tries to deliver everything to the other end, in order and without duplicates. The real time traffic needed to be delivered quickly even though not all of it actually got there and it could even be delivered out of order if necessary. So the split of IP and TCP allowed us to put a base on packet transport including real time traffic. We tried things like voice over IP, which had already been taken into account in the 1970s but didn't become visible to and popular with the public until 1990[7]. But we were already testing it in the 1970s.

7 The first introduction to the public was by a company called Vocal Tech in Israel in 1996. Skype has been a very good example of public interest in or demand for voice over IP. I would estimate today that at the end of 2010 there are possibly around three billion Internet users. This is because many of them use their mobile phones to connect to the Internet.

“What is special about VOIP is that it's just another thing you can do on the Internet, whereas voice is the only thing—or nearly the only thing with the exception of the dial-up modem and fax—that you can do on the public telephone network.”

After working at Stanford, ARPA asked me to go to Washington to manage the Internet research program for them. I accepted and in July or August 1976 I went to Washington and have stayed there (although not at ARPA) since then; it’s now been thirty-one years. I committed myself to running the Internet program, the packet radio program, the satellite packet program and the security program, which I inherited from **Steve Walker,** another ARPA program manager.

I was with ARPA until the end of 1982 and then left the Government to join MCI in order to develop what was then called MCI Mail, which was a commercial email service launched on September 27th, 1983. I stayed after the launch and continued to develop it until June 1986, when Bob Kahn, who had left ARPA at the time, decided to create a company called the Corporation for National Research Initiatives (CNRI). That research organization was formed to focus on information infrastructures and Bob asked me if I wanted to participate in the research. I accepted and in June 1986 I joined CNRI. I stayed there until October 1994, when I was invited back to MCI to introduce them to the Internet business. I agreed to do that and rejoined them as Senior Vice President for Data Architecture. In the end, I inherited the Internet engineering activity at MCI and stayed with the company during its worst moments, because Bernard Ebbers and his organization acquired it in 1998 and in four years it was destroyed. In June 2002, the company filed for bankruptcy and I suppose I could have gone away with a lot of other people, but I thought it was more ethical to stay and help restore it. I don’t claim any important role in that,

but I figured I managed to do all I needed to do for the people who reported to me. I stayed at MCI until the end of 2005, and at this time it was announced that Verizon would acquire the company. To be honest, at the time I thought I had done a lot of work on the Internet infrastructure during the previous eleven years and that it would be reasonable and interesting for me to go back to applications again, just as I had done with Bob Kahn at CNRI. It occurred to me that Google would probably be the most interesting place to go for Internet applications, so I sent an email to Eric Schmidt, who was the president of Google and whom I had known for several years, asking him if he needed any help. He answered in a short note saying yes, and so I joined the company in October 2005. I have been there since then. We discussed what my title should be and when they asked me what title I preferred I suggested "Archduke." It seemed like a really good title but they reminded me that the previous archduke was Franz Ferdinand and that he was assassinated in 1914, an event which triggered off World War I. So I concluded that being an archduke was perhaps not such a good idea. They suggested to me, given the work I had been doing over the past thirty years, that I should be head of Internet evangelism at Google. So, I took on the title of Chief Internet Evangelist, and I spend a lot of time traveling to carry out the role.

In your opinion, what are the key characteristics of the Internet?

I think the most important characteristic is that **it is agnostic in terms of transport.** It doesn't matter whether the packet is transported via fiber optics, satellite channel, radio, fixed line or Ethernet. It just doesn't care and also doesn't know what it is carrying. All it knows is that it is carrying these Internet packets but it doesn't know what the bits in the packets mean. So, content in Internet applications is subject to interpretation by computers at the edges of the network, and not by the network itself. I believe that this principle—the principle of end-to-end connectivity—is vitally important. What it has done is support an explosion of applications that people had been able to write that are relatively insensitive to the basic Internet architecture. It only operates in it, so no permission from an ISP is required to experiment with new ideas. It is simply done, and this

condition generates great incentives for new product development. You don't need anyone's permission to try anything.

What do you consider the most important milestones in the development of the Internet?

- 1972-1973. When Bob Kahn and I started talking about this.
- December 1974. First complete TCP specification. It had mistakes, but Yogen Dalal, Carl Sunshine and I programmed that version. Very quickly thereafter implementation had begun and a series of redesigns had occurred.
- On November 22nd, 1977, we did a three-network demonstration of the Internet using the San Francisco Bay Area Radio Network, the Atlantic Satellite Network and the ARPAnet cable line, which then extended to Norway and University College London.
  So that demonstration showed that you can actually link all three networks together and have them perform successfully.
- On January 1st ,1983, the largest Internet roll-out of the Internet took place to all of the ARPA spon- sored research community which requires to switch away from NCP to use the new TCP/IP protocols suite.
- The first appearance of commercial services takes place in 1989. Although I think the guys at UUnet would disagree and claim that they already had commercial Internet services in 1987. They just didn't connect to NSFnet, which seems right to me. NSFnet allowed the first government-backed commercial service in 1989. Bob Kahn and I agreed to connect MCI Mail to the Internet in 1988 and, in June or July 1989, we launched the MCI Mail interconnection and had it operational by mid-year.
- The ARPAnet network withdrew in 1990.
- The NSFnet which started in 1986 was retired in 1995 and so these milestones demonstrated that the Internet had a life of its own. It has commercial services for those who are able to afford them and that doesn't mean it was necessary to have research components in order to be operational.
- Tim Berners-Lee started the World Wide Web in 1989.

- In 1992, Marc Andreessen and Eric Bina from the National Center for Supercomputing Applications wrote the Mosaic version of the WWW browser. Eric does not always get the kind of visibility he should have, a little bit like Robert Cailliau (CERN) in the case of the World Wide Web.
- 1994 Marc Andreessen and Eric Bina (from NCSA) moved out to Palo Alto with Jim Clark (who started Silicon Graphics and Telemetry) and they started Netscape Communications. By this time, I had returned to MCI and the guy who was with me at the time, Bob Harcharik, and I went to Netscape Communications to see if we could license their customer technology for an e-commerce application we were planning to make at MCI. And, once again, MCI was nearly ten years ahead. We released the MCI Mail with a browser-based interface, but it wasn't very satisfying because no one was really aware of the existence of the Internet yet. But that was an important milestone because it was literally like an online shopping mall, we had people put up web pages and you could go and buy things and complete transactions online. It was not very successful and, in some respects, neither was MCI Mail, which started in 1982 and was operational in 1983, when not very many people were used to being online.
- 1996 was the year of the "dot-com" boom.
- In April 2000, "dot-com" exploded along with the investments and so on. The Internet boom finally fell because many people did not have an adequate business models in place.
- Then came the Internet renaissance during the period comprising 2006 and 2007, when there was a continued demand for Internet services in spite of the decline suffered by several Internet application companies. The demand continued to grow at about 100 percent per year and today I would estimate it to be about 40 percent per year but that is still an enormous growth rate. I think that number is vastly inflated by mobiles that are currently Internet- enabled. There were three billion mobile phones on the market in 2007, of which it is believed that 10% were connected to the Internet. This equates to around 300 million. And if we factor in the other 500 million devices that also had a connection, there were a total of 800 million devices with an Internet connection. So that brings us more or less to the

present where we are struggling with expansion of the Internet, dealing with things like IPv6 and the addition of non Latin domain names into the Internet by ICANN

Vint Cerf preaching in the refectory or dining room of the Cistercian Monastery of Poblet (Tarragona). Photo courtesy of Joan Batet i Pons.

How did you contribute to the development of the ARPAnet?

In terms of my own contribution, I would highlight the original design, in collaboration with Bob Kahn, of the TCP and the subsequent division into TCP/IP. I served on the IAB for many years, and as its chair for a while. I founded the Internet Society with Bob Kahn and others, and was its president for three years. Eventually, I joined the board of ICANN in 1999 and then served as its chairman from November 2000 to November 2007. I was succeeded by Peter Dengate Thrush, who

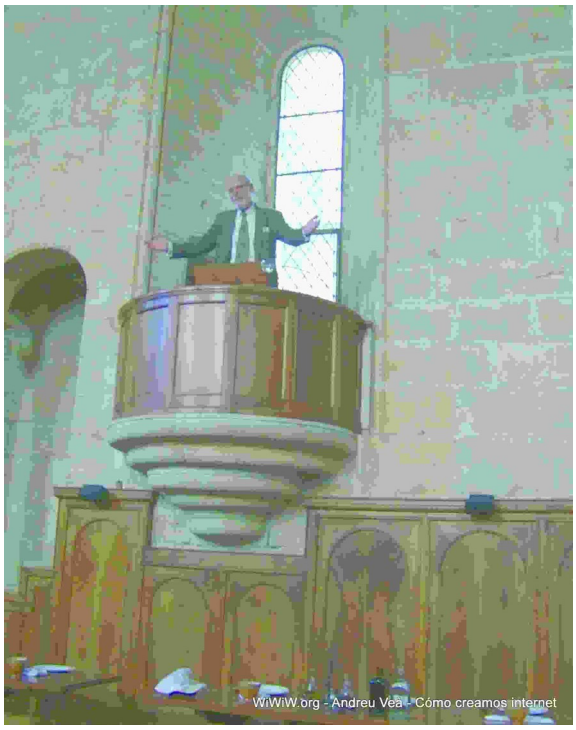

is an intellectual property barrister in New Zealand. So I think that I have made consistent contributions. The functioning of the Internet through interplanetary space was already under way. Work began in 1998 at the Jet Propulsion Laboratory in Pasadena, California. I hope NASA and other international space agencies will adopt our DTN protocols for interplanetary communications, which can address the adverse effects of long and uncertain delays in such communications. The delay could easily be hours when you get all the way to the outer planets. At the speed of light, Mars is 40 minutes away at the most and 6 minutes away at the least, but when we talk about Saturn, Jupiter, Uranus, etc., we are talking about hours of travel and many uncertainties as well.

I think my contributions cover different aspects. At Google, I am trying very hard to work on application side, and trying to help my colleagues at Google here to come up with new ideas and to deal with new ideas in Internet applications.

I have actually been to two of the Internet Governance Forums, one in Athens and most recently in Rio de Janeiro. I missed the most recent one in India. The people in the group and I contributed to that as well, thinking what Internet governance actually means.

Who are some key people in the development of Internet, leaders or trendsetters?

Related to the Internet:

- **Robert E. Kahn** (ARPAnet and Internet); **David Clark; Jon Postel; Robert Braden; Stephen T Kent; Dan Lynch** (head of migration from NCP to TCP on January 1st, 1983 at ARPAnet); **Yogen Dalal, Carl Sunshine, Richard Karp, James Mathis, Ronald Crane,** all of them Stanford graduates with whom I worked at the TCP in my Stanford lab; **Dennis Jennings** from NSF (National Science Foundation) for choosing TCP/IP for NSFnet; **Stephen Wolff** (NSFnet), **Hans-Werner Braun** (Merit-NSFnet); **Gerard LeLann** (worked at Stanford on the design of the TCP although he was from IRIA, France); **John Shoch and Bob Metcalfe,** who worked at

Xerox PARC and attended my Stanford seminars on TCP; **Peter Kirstein** and his students at UCL (University College London); **Ray Tomlinson and Bill Plummer,** from BBN; **Noel Chiappa** (MIT), who worked with *routers;* **Virginia Strazisar** (who designed the first portal in BBN). It is a very long list, especially when considering all the people from 1983 onwards. **Larry Landweber** and **David Farber** (CSNET); **David Crocker, John Vittal, Ray Tomlinson,** for designing one of the first email services for ARPAnet and the Internet.

From the ARPAnet era:

- **Stephen D. Crocker,** for his pioneering work in the predecessor of the TCP: the NCP that worked during the origins of ARPAnet until 1983; **Larry Roberts** (ARPAnet), **Len Kleinrock** and **Howard Frank** (for defining ARPAnet); **Frank Heart, Bob Kahn, Dave Walden, Severe Ornstein, Willy Crowther** and the rest of the BBN team for building and designing the first IMP and network; **Donald W Davies and Roger Scantlebury,** among others, from the National Physical Laboratory in London, the United Kingdom (as inventors of packet switching); **Paul Baran** (RAND, USA), another inventor of packet switching.

Some anecdotes

There is one involving the packet radio network that SRI International was implementing and testing. We did many demonstrations of the packet radio transmission system interconnected to computers on ARPAnet. Whenever I attended these demonstrations, I could tell if the package worked or not because I received the signal through a hearing aid. It transmitted via the 1,710 to 1,850 MHz band and the hearing instrument was sensitive to these frequencies. This meant that you heard the crackling sound of paper whenever the packet was being sent. As soon as the sound stopped I knew that the radio was dead. So, I kicked the guy from SRI under the table for him to restart the radio. This would happen while you were sitting in the conference room with a packet radio system. I only needed to be within a few meters of the radio to detect if the packet was being transmitted, so it was fun.

Vinton Cerf

Vint Cerf with Andreu Veà during the interview on December 8th, 2007, in McLean (Virginia).

What do you think about the future of the Internet?

In terms of the future of the Internet, I am convinced that it will continue to expand, reaching five billion users—maybe even six billion. I believe that this figure will be reached by 2015. There are only eight years to go, but if I'm not mistaken in my prediction of three billion by 2010, this gives us another five years to get the other half, and most of this will be done via mobile phones with Internet access. So it doesn't seem completely impossible.

I also think that we are will be able to enjoy all types of media on the Internet: radio, television, print and so on. It's already happening so it's predictable. I think people are going to use the Internet to control their electronic devices so that all household, office and car appliances and devices or those you carry around with you will be Internet-enabled. I think we are going to have to struggle to figure out how to index and interpret digital content on the network. I am worried about losing our ability to interpret the bits. If you have files for the bits but there are no programs that know how to interpret them, the data and images will lose all their meaning. It's would just be bits, which would be terrible. So it's important

“It's important that we find a way to make sure we can continue to interpret the bits that are on the Internet.”

that we find a way to make sure that we can continue to interpret the bits that are on the Internet, and this probably means that we not only need to update the copies with new applications, but also try to use old ones that would otherwise have disappeared. Someone said: “I’m not using this application anymore.” We, the Internet user community, would like to have access to that software to make sure it still works because, otherwise, our files would no longer be accessible. This raises questions about intellectual property, but I think we have to work this out because otherwise we will end up with what we call "bit rot".

Do you see any technological trends?

The clearest trends today are increased mobility, in- creased bandwidth at the edges of the Internet, in- creased symmetry in network access services rather than

the asymmetric bandwidth that currently exists. I think there will be more mobiles connected and I hope there will be a lot of interplanetary interaction via the Internet.

FURTHER READING

PAPERS & BOOKS MENTIONED/RECOMMENDED

- Cerf, V. and Kahn, R., "A Protocol for Packet Net- work Intercommunication," *IEEE Transactions on Communications* 22, 5 (May 1974), pp. 637-648

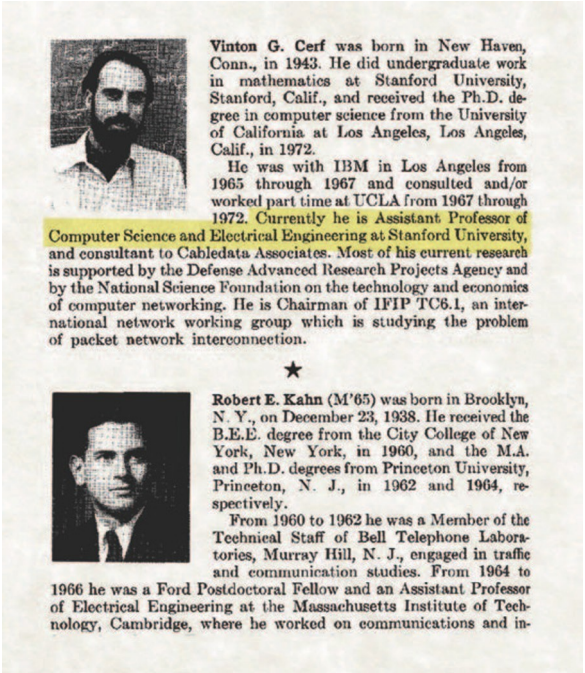

Vinton G. Cerf was born in New Haven, Conn., in 1943. He did undergraduate work in mathematics at Stanford University, Stanford, Calif., and received the Ph.D. degree in computer science from the University of California at Los Angeles, Los Angeles, Calif., in 1972.

He was with IBM in Los Angeles from 1965 through 1967 and consulted and/or worked part time at UCLA from 1967 through 1972. Currently he is Assistant Professor of Computer Science and Electrical Engineering at Stanford University, and consultant to Cabledata Associates. Most of his current research is supported by the Defense Advanced Research Projects Agency and by the National Science Foundation on the technology and economics of computer networking. He is Chairman of IFIP TC6.1, an international network working group which is studying the problem of packet network interconnection.

★

Robert E. Kahn (M'65) was born in Brooklyn, N. Y., on December 23, 1938. He received the B.E.E. degree from the City College of New York, New York, in 1960, and the M.A. and Ph.D. degrees from Princeton University, Princeton, N. J., in 1962 and 1964, respectively.

From 1960 to 1962 he was a Member of the Technical Staff of Bell Telephone Laboratories, Murray Hill, N. J., engaged in traffic and communication studies. From 1964 to 1966 he was a Ford Postdoctoral Fellow and an Assistant Professor of Electrical Engineering at the Massachusetts Institute of Technology, Cambridge, where he worked on communications and in-

Vint Cerf and Bob Kahn as they appear in their famous article "A Protocol for Packet Network Intercommunication" of 1974, in which they define the TCP protocol.